

10 / 1188

1

P/63767.WOP

COMMUNICATIONS SYSTEM

The present invention relates to the field of communications systems in general and to systems for protecting against the effects of equipment failure in communications systems in particular.

- 5 One of the most important concepts in network management is maintaining the survivability of networks. When there are either link or node failures any affected routes should be repaired as soon as possible. A node failure can be considered as the failure of multiple links, that is, a node failure affects the traffic the same way as if all links terminated on it were faulty. The present invention applies to both link and node diversity protection schemes: here diversity
- 10 relates to the property of the resources used by the protection path (links and/or nodes) to be fully disjoint from those used by the working path.

Because of the large volumes of traffic that networks are expected to carry, resulting from the continued explosive growth of data-oriented applications, network survivability has become an issue of paramount importance. In conjunction, there is a continuing drive for maximising

15 efficiency and minimising costs in large networks. Nodes are able to monitor the status of the connections passing through them to detect faults. In the SDH transport network, for example, this can be achieved using both trail monitoring and the so-called non-intrusive monitoring, both defined in the ITU-T specification G.783.

Traditional protection schemes, mainly used in ring networks, consume large amounts of

20 bandwidth. Shared protection (where resources are shared between a number of protection

BEST AVAILABLE COPY

paths) on mesh networks requires less additional capacity to provide fault protection than on ring networks. Although shared protection mesh networks consume fewer network resources, the trade-off has been in longer service restoration times. There is therefore a need for an appropriate architecture to enable fast restoration in such networks.

- 5 The present invention provides a data communications system comprising a plurality of nodes and a plurality of links for providing connections between the nodes; in which a subset of the links and nodes form a worker path for carrying worker data through the communication system; in which the system comprises a further subset of links and nodes for forming a plurality of protection paths for carrying non-worker data in the absence of a fault in the
- 10 worker path and each for providing an alternative path for the worker data in a different part of the worker path in the event of a fault in the worker path; in which the system comprises protection means, in which the alternative paths are predetermined by the protection means prior to the detection of a fault in the worker path; in which the protection means is arranged to activate the entire plurality of protection paths to carry the worker data upon detection of a
- 15 fault in the worker path; in which the protection means is arranged to identify the location of the fault, to return the worker data to those parts of the worker path not affected by the fault and to de-activate any of the protection paths providing an alternative to those parts of the worker not affected by the fault.

According to a preferred embodiment, the system comprises means for allocating the links

20 and nodes one or more cost values relative to the links and nodes of the worker path and

means for selecting on the basis of the one or more cost values a further subset of the nodes and links to form a protection path for at least one link or node of the worker path

The present invention also provides a method of protecting a worker path in a data communications system comprising a plurality of nodes and a plurality of links for providing
5 connections between the nodes; including the steps of passing worker data through a subset of the links and nodes making up the worker path and designating a further subset of links and nodes to form a plurality of protection paths; in which the protection paths carry no worker data in the absence of a fault in the worker path and in which each provides an alternative path for the worker data in a different part of the worker path in the event of a fault in the worker
10 path; including the steps of detecting a fault in the worker path and activating the entire plurality of protection paths to carry the worker data upon detection of a fault in the worker path; including the steps of identifying the location of the fault and returning the worker data to those parts of the worker path not affected by the fault and de-activating any of the protection paths that are providing an alternative for those parts of the worker path not affected
15 by the fault.

According to a preferred embodiment, the present invention also provides a method including the steps of allocating the links and nodes one or more cost values relative to the links and nodes of the worker path and selecting on the basis of the one or more cost values a further subset of the nodes and links to form a protection path for at least one link or node of the
20 worker path..

BEST AVAILABLE COPY

Embodiments of the invention will now be described by way of example with reference to the drawings in which

5

Figures 1 to 3 show conventional operation with protection switching inside a transport network element (TNE);

10

Figures 4 to 7 show a network of TNEs illustrating various protection mechanisms according to the present invention;

Figures 8 and 9 show various aspects of communications systems to which the present invention may be applied.

15

The invention will be explained by way of example with reference to SONET/SDH transport networks due to the great popularity of this kind of network, however the present invention is not limited to such networks. In conventional SONET/SDH networks, one can recover from a faulty link or node in 50 ms or less, but achievement of this depends on optical network links

having a 1+1 backup, i.e. with one backup link allocated for every worker link. To avoid this waste of resources, the present invention provides a shared restoration mechanism that provides enhanced network survivability while minimising the network resources consumed.

- 5 The following table summarises four possible protection schemes with respect to path calculation (i.e. calculation of the sequence of links and nodes to be traversed), resource selection (i.e. selection of the channels to be used within the links) and resource allocation (i.e. the actual implementation of cross connections in the nodes, and consequent seizure of the channels in the links).

<i>Protection Scheme Type</i>	<i>Path Calculation</i>	<i>Resource Selection</i>	<i>Resource Allocation</i>
<i>1</i>	After Failure	After Failure	After Failure
<i>2</i>	Before Failure	After Failure	After Failure
<i>3</i>	Before Failure	Before Failure	After Failure
<i>4</i>	Before Failure	Before Failure	Before Failure

10

The present invention relates to protection scheme type three of the table while conventional SDH protection (APS or Automatic Protection Switching) is type four.

Type three resources are only used after a fault occurs. The same resource can be allocated to several restoration paths, provided that these restoration paths are never activated simultaneously by a single fault. Sharing resources between different restoration paths in this way is the definition of "shared restoration" for this invention.

5

SNCP Operations in a transport network element (TNE)

The cross-connection mechanism needed in a node in order to switch, in a non traffic-affecting way, between the worker and protection path and vice-versa will now be illustrated, with reference to Figures 1 to 3. As stated above, we adopt, by way of example, the SDH
10 nomenclature, but the same conceptual mechanism can be used in all transport network architectures.

Figure 1 illustrates how a single sub-network connection protection (SNCP) protection scheme is implemented, as defined in ITU-T rec. G.805 sec. 7.2.2. Worker ports A and B
15 exchange data via the switch matrix of the TNE. Protection port C is protecting worker port B: if a failure occurs affecting port B, the pre-planned protection starts using port C. Protection uses a split function and a select function within the TNE. The selection module (Sel.) selects the best signal between the two it receives (i.e. one from each of ports B and C) and sends the selected signal to port A via the matrix. The split module splits the signal it receives from port
20 A via the matrix into two signals one sent to each of ports B and C. Worker port A has no

corresponding protection port: a failure affecting part A will adversely affect communications. The relevant information, relating to both the worker (W) and the protection (P) paths, needed in order to perform protection switching, should the need arise is stored in a protection table. The table may be provided in any implementation of a control unit for the TNE.

5

Figure 2, shows a TNE similar to that of Figure 1, with the addition of protection port D. According to Figure 2, worker port A is protected by a SNCP, via protection port D while worker port B is protected by a SNCP via protection port C. An additional selection function (Sel.) selects the best signal between the two it receives from ports A and D and sends it to the
10 "B side" split function via the matrix. An additional split function splits the signal it receives from the "B side" select function via the matrix in two, sending one to each of ports A and D. This behaviour assures that, in case of a failure of one or more resources of the worker path, the traffic is correctly transported by corresponding resources of the protection path. Moreover, it is possible, when the worker path becomes available again, to disable the
15 protection path and revert to the worker path without adversely affecting the transported traffic. A protection table contains the relevant information needed in order to perform protection switching for either port A or B, should the need arise.

Figure 3 depicts how a TNE works in the absence of failures. The TNE of Figure 3 has two
20 worker ports (A, B) and two protection ports (D, C). As the figure shows, during normal

operation no protection scheme is implemented. Worker traffic passes bi-directionally between worker ports A and B via a switch matrix internal to the TNE. A protection table contains the relevant information needed in order to perform protection switching, should the need arise. Activation of the protection table will cause implementation of a double protection
5 scheme as shown in figure 2.

A typical communications network comprises one or more worker paths that cross a set of resources in the network, namely nodes (e.g. TNEs) and links. If only faults occurring among links are considered, the protection is defined as "protection with link diversity". If faulty
10 nodes are considered in addition to links, the protection is defined as "protection with node diversity".

According to the present invention, a preplanned local repair scheme may be configured, activated and deactivated according to the following sequence:

15

1. when the worker path is configured, a set of detours is also computed whose union can bypass all the resources (with either link or node diversity) in the worker path. For shared restoration, details of alternative paths (i.e. the links and nodes used for the detours) are stored for possible future use in the event of a failure but not implemented immediately.
20 Detours must be selected with specific criteria to allow efficient sharing of the protection

resources: see "Selecting and Configuring Detours", below. The alternative paths are not used for carrying the worker data unless a fault has been detected in the worker path.

2. When a failure occurs, locating exactly which resource(s) (link or node) is faulty in the worker path may take a long time, but the source and destination nodes and, possibly, a set of intermediate nodes can very quickly detect that the worker path is faulty somewhere: see "Designating Monitoring Points", below.
3. As the presence of a fault is detected but not located, all the detours are activated at once using SNCP to simultaneously activate multiple detours: see "Activating Detours", below. This achieves the necessary protection in the shortest recovery time.
4. Later, when the fault is fully located by the nodes, only the detour that actually bypasses the faulty resource is left in place: see "Selecting the Protecting Detours", below. If two detours protect the same faulty resource they are both activated, and it is simplest to allow them to remain activated until the fault is fixed, although, one could be deactivated.
5. All the other detours are released, i.e. a hitless partial deactivate is performed: see "Reactivating the Unused Detours" below.

6. When the fault is repaired, the last detour is also released.

Selecting and Configuring Detours

When a path with pre-planned protection is requested, a worker path and a set of protection
5 detours are computed. The worker path may typically be computed by means of a shortest
path algorithm according to conventional techniques.

For each resource (link or node) forming part of the worker path, there must exist at least one
detour to accommodate a fault on that resource. Each detour will preferably use as few
10 resources as possible.

As indicated above, it may happen that two detours protect the same resource. This is
unavoidable if node diversity is required, because, if two adjacent nodes are protected by
different detours, then the link between them will be protected by both these detours. This is
15 illustrated in Figure 4. According to the embodiment of Figure 4, the worker path extends
from source node S to destination node D through nodes 2, 3 and 4. The network of Figure 4
also has three detour paths: detour 10 extends between nodes S, 6 and 3; detour 12 extends
between nodes 2, 7 and 4 and detour 14 extends between nodes 3, 8 and D. Source node S and
destination node D perform trail monitoring. For example nodes 2 and 3 on the worker path
20 are protected by separate detours. (Node 2 is protected by detour 10 and node 3 by detour 12.

As a result, the link between nodes 2 and 3 is protected by both detours 10 and 12). It will be noted that the three detours (10, 12 and 14) of Figure 4 provide a number of disjoint (i.e. not directly connected) alternatives to the worker path. However, for simplicity, such a combination of detours will be referred to here as a single alternative path. This is valid as it is
5 only the combination of all these detours that offer protection for every resource of the worker path.

Two kinds of nodes take part in a detour: nodes that terminate the detour (TD), (i.e. those that are also crossed by the worker path), and nodes that are intermediate in the detour (ID). An ID
10 node stores information defining the cross-connection required for the detour without implementing it unless a fault is detected in a resource of the worker path protected by that ID node, as described below. The ID node may also associate a unique path identifier with the cross-connection definition, in order to allow it to be identified when activation of the detour is requested (see below).

15

In normal, fault-free operation, a TD node implements the "worker" cross-connection that forms part of the worker path. The TD node stores information defining a "detour" cross-connection i.e. a "fork" in the path that would transfer traffic from the worker path to the detour, i.e. the protection path.

20

Designating Monitoring Points

The speed at which detours are activated is improved according to a preferred embodiment by sending the ActivateDetour message from both the beginning and end of each detour rather than from just one end (the ActivateDetour message contains the unique path identifier for the cross-connection definition). This reduces the time required for signaling to the propagation delay across half of the longest detour. The signaling will start at a monitoring TD node when that node detects a failure in the worker circuit. For these reasons, placing non-intrusive monitoring functions at all TD nodes is recommended. Examples are given in the drawings where Figure 5 shows an example in which not all nodes in the worker path implement monitoring points (i.e. only nodes 4 and 7, see below) and Figure 6 shows an example in which all nodes in the worker path implement monitoring points

According to the embodiment of Figure 5, the worker path extends from source node S to destination node D through nodes 2, 4, 5, 7 and 8. The network of Figure 5 also has three detour paths that extend, respectively between nodes S, 3 and 4; 4, 6 and 7 and 7, 9 and D. Source node S and destination node D perform trail monitoring. Only nodes 4 and 7 perform non-intrusive monitoring. According to the embodiment of Figure 6, the worker path extends from source node S to destination node D through nodes S, 2, 3, 4, and D. The network of Figure 6 also has three detour paths that extend, respectively between nodes S, 6 and 3; 2, 7 and 4; and 3, 8 and D. Source node S and destination node D perform trail monitoring. Nodes 2, 3 and 4 perform non-intrusive monitoring

Activating Detours

When the (TD) nodes implementing monitoring detect a worker path failure, they all send an ActivateDetour message through the detour paths. The ActivateDetour message contains a
5 unique path identifier to inform the ID nodes which detour connections must be activated. The ActivateDetour message actually identifies to the ID nodes which parts of the detour path to activate (i.e. by use of the unique path identifier for the cross-connection definition referred to above under "Selecting and Configuring Detours"). We may think of a working path and all its associated detours as a single object with a unique identifier, if any node receives an
10 ActivateDetour message with a certain identifier, it activates the connections implementing the detours associated the that identifier. This includes the TD nodes implementing the appropriate detour cross-connections (SNCP forks).

The ID nodes, in turn, propagate a received ActivateDetour message unless they have already
15 received it, (i.e. form the TD node at the other end of the detour connection) and implement the previously defined cross-connection corresponding to the unique path identifier in the message.

If a node is trying to propagate an ActivateDetour message to a faulty resource (i.e., a link
20 identified as faulty by a disabling alarm: a faulty node is perceived as a faulty link by its

neighbouring node), the cross-connection through the node leading to the faulty resource is deleted locally and autonomously by the node itself and a DeleteDetour message is generated by the node to recover that unusable detour (see "De-activate the unused detours", below).

5 Selecting the Protecting Detours

A single fault on a specific resource will be protected by one or more detours. It is possible that further detours are present in the network that are not protecting the faulty resource and can therefore be released, i.e. all the detours are implemented at first and then some de-activated. When the faulty resource is identified and its location determined (e.g. by node
10 diagnostics), it is possible to identify which detours are protecting the faulty resource and which ones not. The latter can then be released to allow them to carry other traffic and to protect other faults, if any, thus enhancing network resilience. This can be viewed as an early partial-deactivate: i.e. the de-activation of some detours before the fault is repaired.

15 The identification of non-protecting detours will now be described with reference to Figure 7. Figure 7 shows a network with overlapping detours similar to that of Figure 6. According to the embodiment of Figure 7, the worker path extends from source node S to destination node D through nodes 2, 3, 4 and 5. The network of figure 7 also has three detour paths that extend, respectively between nodes S, 8 and 3; 2, 7 and 5; and 3, 9 and D. The network of
20 Figure 7 has a faulty resource (see spike symbol) i.e. the link between TD nodes 3 and 4 on

the worker path. The two detour paths crossing nodes 7 and 9 are both protecting the faulty resource, while the detour crossing nodes 8 is not and can therefore be released. To locate the non-protecting detours the following algorithm is used:

- 5 - (1) indicate with $\langle L$ and $L \rangle$ the nodes where the alarm of the faulty resource has been detected as follows: indicate nodes where a fault has been detected towards the destination with $\langle L$ and nodes where a fault has been detected towards the source of the worker path with $L \rangle$. $\langle L$ and $L \rangle$ may be thought of as status fields within the nodes.

- 10 - (2) indicate with $\langle P$ nodes where a detour begins in the direction of the destination. Indicate with $P \rangle$ nodes where a detour begins in the direction of the source ($\langle P$ and $P \rangle$ are only present in TD nodes). $\langle P$ and $P \rangle$ may also be thought of as status fields within the nodes. $\langle P, P \rangle$, $\langle L$ and $L \rangle$ are not mutually exclusive: see figure 7.

- 15 Hence, in the embodiment of Figure 7, the TD nodes have the following status:

S	=	$\langle P$;	4	=	$L \rangle$;
2	=	$\langle P$;	5	=	$P \rangle$;
3	=	$\langle P \quad P \rangle \quad \langle L$;	D	=	$P \rangle$.

According to the embodiment of Fig. 7, only nodes 3&4 detect the fault, although in practice more of the TD nodes could do this.

- 5 - (3) nodes with <L send a RevertSource message along the worker path in the direction away from the faulty resource, i.e. towards the source node.
- 10 - (4) if a node with a <P receives the RevertSource message, the message is propagated along the worker path and no action is taken. It is possible for the same node to be indicated with both <P and <L. If so, the algorithm works as if there were a message passed between a first node with <L and a second node with <P, although, in practice, there is no need of a message
- 15 - (5) if a node indicated by a P> receives the RevertSource message, the message is propagated along the worker path and the detour is recovered as described below. It is possible for the same node to be indicated with both P> and <L. If so, the algorithm acts as if a message is passed from <L to P> in a similar way to the case described at step (3), above for <P and <L,

- (6) the node with $L>$ sends a RevertDestination message along the worker path in the direction away from the faulty resource, i.e. towards the Destination node.
- (7) if a node indicated by $P>$ receives the RevertDestination message, the message is
5 propagated along the worker path and no action is taken. It is possible for the same node to be indicated with both $P>$ and $L>$. If so, the algorithm acts as if a message is passed from $L>$ to $P>$ in a similar way to the case described at step (3), above for $<P$ and $<L$,
- (8) if a node indicated by $<P$ receives the RevertDestination message, the message is
10 propagated along the worker path and the detour is recovered as described below. It is possible for the same node to be indicated with both $<P$ and $L>$. If so, the algorithm acts as if a message is passed from $L>$ to $<P$ in a similar way to the case described at step (3), above for $<P$ and $<L$.
- 15 Turning to the embodiment of Figure 7, a RevertDestination message is sent from node 4 ($L>$) and propagated through node 5 ($P>$) to destination node D ($P>$): nothing is done, because the message crosses only nodes with $P>$ representing protecting detours.

Simultaneously, in Figure 7, a RevertSource message is sent from node 3 ($<P$, $P>$ and $<L$) and
20 propagated through node 2 ($<P$) to source node S ($<P$). When the indications $<P$ are

encountered in node 3 and then in nodes 2 and S, nothing happens. When the indication P> is encountered in node 3, the detour via node 8 is deactivated and the resources recovered.

Deactivating the Unused Detours

- 5 De-activation of an unused detour is always initiated in a TD node. The worker signal on the detour is forced onto the worker path, then a DeleteDetour message is sent along the detour path from the TD node initiating detour deactivation and all the detour cross-connections are deleted. Information relating to the detour cross-connection in TD and ID nodes continues to be maintained, together with its associated path identifier, for further use in the event of future
- 10 faults.

The process of selecting the protection path will now be described by way of example with reference to Figures 8 and 9 and using the following notation and definitions.

15 Notation:

\exists : existential quantifier that reads "there exists".

$\{a, b, c\}$ the set containing a, b and c;

\in set membership;

\subseteq subset;

20 \cup union;

\cap intersection;

\neg negation;

\emptyset empty set

5 Definitions:

Let the network be described by a graph $G = \{N, E, c\}$, where

$N = \{ \dots n_i \dots \}$ is the set of nodes,

$E \subseteq N \times N$ is the set of links, where each link provides a connection between two nodes of the set N . It is the resource that gives bandwidth between two adjacent nodes and is identified by those two nodes;

$c: E \rightarrow \mathcal{R}$ is the cost function

where \mathcal{R} : the field of the real numbers. The cost c is a function that maps links into real numbers; several cost values may be assigned to a link, all expressible with a real number, the main values include financial cost, total capacity, used capacity and remaining capacity; here we use also two auxiliary modified cost functions. In other cases to which the present invention has application (for example in all optical networks), you may need a very complex cost function expressing physical fiber parameters and optical signal quality. For the purposes of the present example embodiment, we assume a simple real cost.

e_f is a faulty link;

n_f is a faulty node

A path wp on the network is described by a sequence of nodes and their corresponding links $wp = [n_1, e_1, n_2, e_2, \dots, n_{k-1}, e_{k-1}, n_k]$ provided link e_i connects nodes n_i and n_{i+1} , i.e. it satisfies the property $e_i = (n_i, n_{i+1})$. We will call n_1 the source of wp and n_k its destination.

Given a path wp and a resource (link e_f or node n_f) on the path that it is desired to protect with
 5 a protection path or detour, we define two more cost functions cl and cn as follows:

$$cl(wp, e_f, e_i) = \begin{cases} \infty & \text{if } f = i \\ 0 & \text{if } f \neq i \text{ and } e_i \in wp \\ c(e_i) & \text{otherwise} \end{cases}$$

The cost cl associated with each faulty link (e_f) is infinite (i.e. never use). The cost associated with a working link on the path (wp) is zero (i.e. always use). Otherwise, the cost is given by the function c .

10

$$cn(wp, n_f, e_i) = \begin{cases} \infty & \text{if } \exists n_l \text{ such that } e_i = (n_f, n_l) \text{ or } e_i = (n_l, n_f) \\ 0 & \text{if } (\neg \exists n_l \text{ such that } e_i = (n_f, n_l) \text{ or } e_i = (n_l, n_f)) \text{ and } e_i \in wp \\ c(e_i) & \text{otherwise} \end{cases}$$

The cost cn associated with each faulty node (n_f) is infinite (i.e. never use). The cost associated with a working node on the path (wp) is zero (i.e. always use). Otherwise, the cost is given by the function c . The cost function c is part of the definition of the input network and is taken as
 15 predetermined, e.g. on network initialization.

These two cost functions (cl , cn) are used to find a minimum-cost detour protecting a potential faulty resource (e_f) or (n_f) of a path wp . Cost function cl allocates cost ∞ to link e_f to prevent its use as a detour for itself. Cost function cl allocates cost 0 to links in wp that are already
 20 configured to ensure that these are always selected, where available. Cost function cn allocates costs relating to node n_f in a similar way.

These cost functions operate to favour locality, i.e. resources closest to the resource to be protected. In this way the length of the detour, the amount of resources used and the time taken for implementing the detour are all minimized.

5

The concept of Shared Risk Link Group (SRLG) will now be discussed with reference to the drawings in which Figure 1 shows a path passing through a number of resources. A SRLG indicates a pool of resources that share the same failure risk e.g. the fiber that share the same duct, links that share the same node will belong to the same SRLG.

- 10 A link can belong to more than one SRLG. In figure 1, for example, the two ducts define two SRLGs and the fiber panel defines another SRLG, thus the link from A to B belongs to three SRLGs.

The following naming conventions are used here.

A path between two points A and B will be indicated by AB.

- 15 A contiguous subset of resources in AB is called a segment of AB and is indicated with an index (e.g., AB_1 , AB_i , AB_k).

If e is a link, we will write $SRLG(e)$ for the set of SRLGs that e belongs to.

If path $wp = [n_1, e_1, n_2, e_2, \dots, n_{k-1}, e_{k-1}, n_k]$, we will write $SRLG(wp)$ for $\cup_i SRLG(e_i)$.

A link interferes with another if the two links belong in the same SRLG.

- 20 If g is an SRLG and $g \in SRLG(wp)$, its elements are links.

We will call closure of a path wp (more in general, closure of a link set) with respect to interference of the set of links $IntClos(wp) = \cup_i g_i$ for all $g_i \in SRLG(wp)$

Figure 2 illustrates by way of example the concept of closure of a link set with respect to interference. With respect to figure 2 we can assert the following statements:

AB and CD belong to SRLG2;

$SRLG(AB) = \{SRLG1, SRLG2\}$;

5 $SRLG(\{CD, EF\}) = SRLG(CD) \cup SRLG(EF) = \{SRLG2, SRLG3, SRLG4\}$;

AB interferes with GH and CD;

$IntClos(\{CD, EF\}) = \{AB, CD, EF, IJ\}$.

10 When a path with protection is requested from A to B a working path AB and a set of protection detours $d_1(AB), \dots, d_k(AB)$ should be computed.

By definition, a detour $d_i(AB)$ protects a segment AB_i of the working path.

The working path is typically computed by means of a shortest path algorithm. The protection detours should satisfy the following properties:

15 For each considered resource (link or node) supporting the working path, there should exist at least one detour restoring a fault on that resource: $\cup_i AB_i = AB$. This means that the combination of all detours will provide a complete alternative path between A and B. Advantageously, each detour should use as few resources as possible to allow maximum resource recovery. By setting cost to 0 for the already used resources and to ∞ for the faulty resource,
20 the detour is encouraged to follow AB as closely as possible and to deviate only from the faulty resources. This is achieved by applying a minimal cost algorithm to these costs.

The allowed resources for detour implementation are:

- the free (unused) resources
- the resources already used for detours of other working paths, provided that the segments protected by the two detours have no common point of failure: that is, for
5 $d_i(AB)$ some resources already used for $d_j(XY)$ if AB_i and XY_j have no common point of failure.

The algorithm guarantees the following:

10 $SRLG(AB) \cap SRLG(d_i(AB)) = \emptyset$, for each detour $d_i(AB)$

Otherwise, a single failure could affect working and protection simultaneously.

If $SRLG(AB_i) \cap SRLG(XY_j) \neq \emptyset$, $d_i(AB)$ and $d_j(XY)$ use disjoint resource sets and can therefore be implemented simultaneously.

This is necessary as, if $SRLG(AB_i) \cap SRLG(XY_j) \neq \emptyset$, there exist at least one resource whose
15 failure would affect simultaneously the segments AB_i and XY_j , thus, for effective restoration, it should be possible to allocate simultaneously $d_i(AB)$ and $d_j(XY)$.

Two routes or segments, detours or whatever set of resources are in link diversity if they are
20 not supported by the same link or links, that is there is no single link fault that could simultaneously puts the two sets of resources out of order.

The investigation of resources for protection is done iterating on all links of a path, i.e. considering one link at a time and investigating how best to protect a fault on it.

5 *INIT*

Compute $AB = [A=n_1, e_1, n_2, e_2, \dots, n_{k-1}, e_{k-1}, B=n_k]$ on the free resources in E using cost function c with the algorithm of choice

Initialize the detour set $DS = \emptyset$

10 *Initialize the iterating faulty link $fl = e_1$*

Set $i = 1$

Set all resources used by some working path "busy" and all resource used by detours "free"

ITERATION

15 *Compute AB' on the free resources in E using cost function $cl(AB, fl)$; where AB' is another path from A to B that follows a different route from path AB.*

Set $AB_i = AB - (AB \cap AB')$

Set $d_i(AB) = AB' - (AB \cap AB')$

20 *Find $IntClos(AB_i)$*

Find all $d_j(XY)$ such that XY_j uses some link in $IntClos(AB_i)$

If there are free resources in $d_i(AB)$ not used by any $d_j(XY)$, then

Insert $d_i(AB)$ in DS

If the B side of $AB \cap AB'$ includes only node B, terminate

5 *Set fl to the first link on the B side of $AB \cap AB'$*

Set again all resources used by detours "free"

Increment i

Continue iteration

Else

10 *Set resources used by any $d_j(XY)$ "busy" (temporarily for this iteration)*

Continue iteration

We now consider the algorithm for the protection path calculation with node diversity.

15 Two routes or segments, detours or whatever set of resources are in node diversity if they are not supported by the same node, that is there is no single node fault that could simultaneously put the two sets of resources out of order.

INIT

Compute $AB = [A=n_1, e_1, n_2, e_2, \dots, n_{k-1}, e_{k-1}, B=n_k]$ on the free resources in E using cost function c with the algorithm of choice

Initialize the detour set $DS = \emptyset$

Initialize the iterating faulty node $fn = n_2$

5 *Set $i = 1$*

Set all resources used by some working path "busy" and all resource used by detours "free"

ITERATION

Compute AB' on the free resources in E using cost function $cn(AB, fn)$

10 *Set $AB_i = AB - (AB \cap AB')$*

Set $d_i(AB) = AB' - (AB \cap AB')$

Find $IntClos(AB_i)$

Find all $d_j(XY)$ such that XY_j uses some link in $IntClos(AB_i)$

If there are free resources in $d_i(AB)$ not used by any $d_j(XY)$, then

15 *Insert $d_i(AB)$ in DS*

Set fn to the first node on the B side of $AB \cap AB'$

If $fn = B$, terminate

Set again all resources used by detours "free"

Increment i

Continue iteration

Else

Set resources used by any $d(XY)$ "busy" (temporarily for this iteration)

Continue iteration

5

The invention has been explained above by way of example mainly with reference to SDH Transport Networks due to the great popularity of this kind of network. However, it will be apparent to the skilled reader that the invention can be applied to all forms of transport
10 network and to all topologies including, but not limited to, mesh and ring topologies.

The above-described embodiments are presented by way of example only and are not intended to limit the scope of the invention. The present invention advantageously provides a means and method for fast switching to a pre-determined protection path and a routing algorithm,
15 method and means suitable for computing the detours for a protection path in a shared local repair scheme as well as other forms of data communications systems. This algorithm takes into account interference of working paths to allow repair of multiple services against failure of multiple resources in shared risk group. It also considers maximization of locality to allow the use of the minimum amount of resources during a fault and to speed up detour activation.

20